arXiv:1604.01250v2 [stat.ML] 13 May 2016

ROYAL SOCIETY OPEN SCIENCE

rsos.royalsocietypublishing.org

Research

Article submitted to journal

Subject Areas:
Statistics

Keywords:
Gaussian processes, Regression, Data analysis, Inference

Author for correspondence:
Christopher J. Moore
e-mail: cjm96@ast.cam.ac.uk

Fast methods for training Gaussian processes on large data sets

C. J. Moore[1], A. J. K. Chua[1], C. P. L. Berry[2] and J. R. Gair[3]

[1]Institute of Astronomy, Madingley Road, Cambridge, CB3 0HA, UK
[2]School of Physics & Astronomy, University of Birmingham, Birmingham, B15 2TT, UK
[3]School of Mathematics, University of Edinburgh and Biomathematics and Statistics Scotland, James Clerk Maxwell Building, Peter Guthrie Tait Road, Edinburgh EH9 3FD, UK

Gaussian process regression (GPR) is a non-parametric Bayesian technique for interpolating or fitting data. The main barrier to further uptake of this powerful tool rests in the computational costs associated with the matrices which arise when dealing with large data sets. Here, we derive some simple results which we have found useful for speeding up the learning stage in the GPR algorithm, and especially for performing Bayesian model comparison between different covariance functions. We apply our techniques to both synthetic and real data and quantify the speed-up relative to using nested sampling to numerically evaluate model evidences.

1. Introduction

A wide range of commonly occurring inference problems can be fruitfully tackled using Bayesian methods. A particular common inference problem is that of regression; determining the relationship of a control variable x to an output variable y given a set of measurements of $\{y_i\}$ at points $\{x_i\}$. The solution requires a model $y = f(x)$, which allows us to predict the value of y at an untested value of x. From a Bayesian standpoint, this can be achieved using Gaussian processes (GPs): a GP is collection of random variables, of which any finite subset have a joint Gaussian probability distribution [1].

© 2015 The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License http://creativecommons.org/licenses/by/4.0/, which permits unrestricted use, provided the original author and source are credited.

THE ROYAL SOCIETY PUBLISHING

Gaussian process regression (GPR) is a powerful mathematical technique for performing non-parametric regression in a Bayesian framework [1–5]. The key assumption underpinning the method is that the observed data set being interpolated is a realisation of a GP with a particular covariance function. This assumption presents us with a challenge: how do we choose the covariance function which gives the best interpolant?

The process of choosing the covariance function is known as *learning*, or *training* of the GP. In this training process, it is necessary to compute the inverse of the covariance matrix (the matrix formed by evaluating the covariance function pairwise between all n observed points). The time taken to evaluate the inverse of the covariance matrix scales as $\mathcal{O}(n^3)$ [1], where n is the number of points being interpolated; this has typically restricted the application of GPR to smaller problems ($n \lesssim 10^5$), although work has been done on extending its applicability to larger data sets [6–9].

In this paper, we present two techniques that speed up the training stage of the GPR algorithm. The first aims to reduce the dimensionality of the problem, and hence speed up the learning of the hyperparameters for a single covariance function; this does not change the fact that the cost of this process is $\mathcal{O}(n^3)$, instead it simply reduces the constant in this scaling. The second aims to enable fast Bayesian model comparison between different covariance functions while also incorporating the benefits of the first technique.

We consider maximising the hyperlikelihood: the conditional probability of the data given a particular set of hyperparameters used to specify the covariance function.[1] We provide an expression for the Hessian matrix of the hyperlikelihood surface and show how this can be used as a valuable tool for comparing the performance of two different covariance functions. We also present modified expressions for the hyperlikelihood, its gradient and its Hessian matrix, which have all been analytically maximised and marginalised over a single scale hyperparameter. This analytic maximisation or marginalisation reduces the dimensionality of the subsequent optimisation problem and hence further speeds up the training and comparison of GPs.

These techniques are useful when attempting to rapidly fit large, irregularly sampled datasets with a variety of covariance function models. The authors have previously made use of these techniques in exploring the correlation structure of the differences between complicated waveform models in the field of gravitational-wave astronomy [10, 11]; this was done so that the effect of different models on the parameter infererences could be marginalised over. There, the behaviour of the data was largely unknown a priori and it was necessary to quantitatively compare a wide range of different covariance functions. Work in this area with larger datasets is ongoing.

In Sec. 2, we review the GPR method and discuss methods of efficiently determining a covariance function. In Sec. 2(a), we present our expression for the Hessian of the hyperlikelihood along with a discussion of how it can be used for model comparison, and in Sec. 2(b) we show how the training of the GP can be accelerated by analytically maximising or marginalising the hyperlikelihood over a single scale parameter. In Sec. 3, we apply these methods to both synthetic and real data sets, and compare the computational cost to that of a full numerical evaluation of the Bayesian model evidences. Finally, a brief discussion and concluding remarks are given in Sec. 4.

2. Gaussian process regression & training

The technique of GPR is a method for interpolating (or extrapolating) the data contained in a *training set* $\mathcal{D} = \{\boldsymbol{x}, \boldsymbol{y}\}$. The vector $\boldsymbol{x} = \{x_i | i = 1, 2, \ldots, n\}$ is called the input vector and the output vector is given by $y_i = f(x_i)$ for some unknown function f. The method works by assuming that the data have been drawn from an underlying GP $f(x) \sim \mathcal{GP}(\mu(x), k(x, x'))$ with specified mean $\mu(x)$ (usually assumed to be zero) and covariance function $k(x, x')$.

[1]This is often referred to as the marginal likelihood. In order to avoid confusion with distributions over the model parameters, we prefer to consistently use the hyper- prefix to denote probability distributions connected to the inference of the hyperparameters.

There is freedom in specifying the covariance function; common choices, such as the squared exponential and Matérn function, include a number m of free hyperparameters $\boldsymbol{\theta} = \{\theta_i | i = 1, 2, \ldots, m\}$ that control the properties of the GP, i.e. $k(x, x') = k(x, x'; \boldsymbol{\theta})$.

The predictive power of the method comes from computing the conditional probability of the function taking a given value at some new $(n+1)$-th input point x_*, given the observed values in $\mathcal{D}$ and the values of the hyperparameters $\boldsymbol{\theta}$. This predictive probability distribution $P(y(x_*)|\mathcal{D}, \boldsymbol{\theta})$ for the function at the new point is a Gaussian with mean $\overline{y}(x_*)$ and variance $[\sigma_y(x_*)]^2$ [1],

$$\overline{y}(x_*) = \boldsymbol{k}_*^{\mathrm{T}}\, \mathbf{K}^{-1}\, \boldsymbol{y}\,, \quad [\sigma_y(x_*)]^2 = k_{**} - \boldsymbol{k}_*^{\mathrm{T}}\, \mathbf{K}^{-1}\, \boldsymbol{k}_*\,, \tag{2.1}$$

$$y(x_*)|\{\mathcal{D}, \boldsymbol{\theta}\} \sim \mathcal{N}\left(\overline{y}\left(x_*\right), \sigma_y\left(x_*\right)\right)\,, \tag{2.2}$$

where we have defined the scalar, vector and matrix shorthand

$$k_{**} \equiv k(x_*, x_*), \quad [\boldsymbol{k}_*]_i \equiv k(x_*, x_i)\,, \quad [\mathbf{K}]_{ij} \equiv k(x_i, x_j)\,. \tag{2.3}$$

Since the posterior distribution for (2.2) relies upon the form of the covariance, GPR cannot be used to make definite predictions until we have fixed a method for dealing with the unknown hyperparameters $\boldsymbol{\theta}$.

Ideally, we would like to place a prior probability distribution on $\boldsymbol{\theta}$ and make predictions by evaluating the integral

$$P(y(x_*)|\mathcal{D}) = \int \mathrm{d}\boldsymbol{\theta}\; P(y(x_*)|\mathcal{D}, \boldsymbol{\theta}) P(\boldsymbol{\theta}|\mathcal{D}) = \int \mathrm{d}\boldsymbol{\theta}\; P(y(x_*)|\mathcal{D}, \boldsymbol{\theta}) P(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta}) P(\boldsymbol{\theta})\,, \tag{2.4}$$

where we have used Bayes' theorem to obtain the second equality. We have introduced the hyperlikelihood given by

$$\ln P(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta}) = -\frac{1}{2}\left[\boldsymbol{y}^{\mathrm{T}}\, \mathbf{K}^{-1}\, \boldsymbol{y} + \ln(\det \mathbf{K}) + n \ln(2\pi)\right]\,, \tag{2.5}$$

which encodes the probability that the observed (training) data were drawn from a GP with covariance function k. The integral (2.4) is almost always analytically intractable and prohibitively expensive to evaluate numerically. A common approximate approach is to use the most probable values of the hyperparameters $\hat{\boldsymbol{\theta}}$, which maximise $P(\boldsymbol{\theta}|\mathcal{D})$ [12–14].

Assuming the prior distribution is sufficiently flat (or uninformative) over the region of interest, this is equivalent to maximising the hyperlikelihood $P(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta})$. Under this approximation, the predictive distribution becomes

$$P(y(x_*)|\mathcal{D}) \simeq P(y(x_*)|\mathcal{D}, \hat{\boldsymbol{\theta}})\,, \tag{2.6}$$

which is simply the Gaussian in (2.2) with mean and variance evaluated at $\hat{\boldsymbol{\theta}}$. Implementing the above procedure requires numerically maximising the hyperlikelihood in (2.5). This can be computationally expensive; in Sec. 2(a) and Sec. 2(b), we present methods for reducing the cost of maximising the hyperlikelihood.

(a) Using the gradient & Hessian

The maximisation process may be accelerated if the gradient of the hyperlikelihood is known and a gradient-based algorithm, such as a conjugate gradient method [13, 15], can be used. The gradient of the logarithm of the hyperlikelihood is given by [1]

$$\partial_{\boldsymbol{\theta}} \ln P(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta}) = \frac{1}{2}\boldsymbol{y}^{\mathrm{T}}\, \mathbf{K}^{-1} \cdot \partial_{\boldsymbol{\theta}}\mathbf{K} \cdot \mathbf{K}^{-1}\, \boldsymbol{y} - \frac{1}{2}\,\mathrm{Tr}(\mathbf{K}^{-1} \cdot \partial_{\boldsymbol{\theta}}\mathbf{K})\,. \tag{2.7}$$

This can be shown by differentiating (2.5) and making use of the standard results

$$\partial \mathbf{K}^{-1} = -\mathbf{K}^{-1} \cdot \partial\mathbf{K} \cdot \mathbf{K}^{-1}\,, \quad \partial\left(\det \mathbf{K}\right) = (\det \mathbf{K})\,\mathrm{Tr}(\mathbf{K}^{-1} \cdot \partial\mathbf{K})\,. \tag{2.8}$$

The gradient in (2.7) is useful because the rate-determining step in computing the hyperlikelihood is computing the inverse matrix $\mathbf{K}^{-1}$ (usually achieved through a Cholesky decomposition in

practice), which is an $\mathcal{O}(n^3)$ operation. All other steps in (2.5) scale as $\mathcal{O}(n^2)$ or less.[2] Once the inverse has been calculated, the gradient in (2.7) may also be evaluated in $\mathcal{O}(n^2)$; so in evaluating the hyperlikelihood for a large training set we can also get the gradient for negligible extra cost.

The procedure outlined above can be performed for multiple covariance functions, each yielding a different GP interpolant. It is therefore necessary to have a method of comparing the performance of different interpolants to decide which to use. One way to achieve this is to evaluate the (hyperprior-weighted) volume under the hyperlikelihood surface, the hyperevidence, and use this as a figure of merit for the performance. Evaluating this integral is prohibitive, so an approximation is to calculate the Hessian matrix of the $\ln P(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta})$ surface at the peak (the position and value of which have already been found) and to analytically integrate the resulting Gaussian. This procedure assumes flat (or slowly varying) hyperpriors in the vicinity of the peak, but this has already been assumed in going from (2.4) to (2.6). Differentiating the gradient in (2.7), again making use of the results in (2.8), and evaluating the derivatives at the position of peak hyperlikelihood, $\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}$, gives the Hessian,

$$\begin{aligned}\partial_{\boldsymbol{\theta}}\partial_{\boldsymbol{\theta}'}\ln P(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta})|_{\hat{\boldsymbol{\theta}}} = &-\frac{1}{2}\boldsymbol{y}^{\mathrm{T}}\left[2\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\theta}}\mathbf{K}\cdot\mathbf{K}^{-1}\partial_{\boldsymbol{\theta}'}\mathbf{K}\cdot\mathbf{K}^{-1}-\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\theta}}\partial_{\boldsymbol{\theta}'}\mathbf{K}\cdot\mathbf{K}^{-1}\right]\boldsymbol{y}\\ &+\frac{1}{2}\,\mathrm{Tr}\left(\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\theta}}\mathbf{K}\cdot\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\theta}'}\mathbf{K}-\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\theta}}\partial_{\boldsymbol{\theta}'}\mathbf{K}\right)=-\mathbf{H}\,. \qquad (2.9)\end{aligned}$$

This expression has the same advantages as the expression for the gradient; as the inverse of the covariance matrix has already been computed, the Hessian may be evaluated at negligible extra cost. The hyperlikelihood surface may therefore be approximated by the Gaussian [12, 16]

$$\ln P(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta})\approx\ln P(\boldsymbol{y}|\boldsymbol{x},\hat{\boldsymbol{\theta}})-\frac{1}{2}\Delta\boldsymbol{\theta}^{\mathrm{T}}\cdot\mathbf{H}\cdot\Delta\boldsymbol{\theta}\,. \qquad (2.10)$$

We seek the evidence, which is given by the following integral of the posterior, where we have specified a prior $\Pi(\boldsymbol{\theta})$ on the hyperparameters;

$$\mathcal{Z}(\mathcal{D})=\int \mathrm{d}\boldsymbol{\theta}\;\Pi(\boldsymbol{\theta})P(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta})\,. \qquad (2.11)$$

Assuming the posterior is a sufficiently well peaked distribution, with peak at position $\boldsymbol{\theta}=\tilde{\boldsymbol{\theta}}$, the evidence may be written using the Laplace approximation [2] as

$$\mathcal{Z}(\mathcal{D})\approx\Pi(\tilde{\boldsymbol{\theta}})P(\boldsymbol{y}|\boldsymbol{x},\tilde{\boldsymbol{\theta}})\sqrt{\frac{(2\pi)^m}{\det\left(\boldsymbol{H}+\boldsymbol{H}_{\Pi}\right)}}\,. \qquad (2.12)$$

It is always possible to change the hyperparameterisation so that the prior is flat in which case the hyperposterior is proportional to the hyperlikelihood.[3] If such a hyperparameterisation has been chosen then $\Pi(\tilde{\boldsymbol{\theta}})=1/V$ (where V is the hyperprior volume, or range of integration), $\boldsymbol{H}_{\Pi}=0$, and $\tilde{\boldsymbol{\theta}}=\hat{\boldsymbol{\theta}}$; therefore

$$\mathcal{Z}(\mathcal{D})\approx\frac{P(\boldsymbol{y}|\boldsymbol{x},\hat{\boldsymbol{\theta}})}{V}\sqrt{\frac{(2\pi)^m}{\det\mathbf{H}}}\,. \qquad (2.13)$$

This expression is now invariant under further changes to the hyperparameter specification which preserve the property that the prior is constant. We use hyperparameterisations with flat hyperpriors as this choice uniquely specifies the approximation in Eq. (2.13); although there remains the possibility that another hyperparameterisation exists in which the posterior is better approximated as a Gaussian.

For two covariance functions, k_1 and k_2, the odds ratio may be defined as the ratio of the value of (2.13) evaluated with k_1 to the value evaluated using k_2, and this may be used to discriminate among competing models. The hyperprior volume V in (2.13) acts as an Occam

[2] As described in a footnote in [1], the matrix–matrix products in (2.7) should not be evaluated directly, as this is an $\mathcal{O}(n^3)$ operation. Rather, the first term should be evaluated in terms of matrix–vector products, and, in the second term, only the diagonal elements that contribute to the trace need to be calculated; these are both $\mathcal{O}(n^2)$ operations.

[3] For example, if the original prior on the parameters $\boldsymbol{\theta}$ is $p(\boldsymbol{\theta})$, we can define $\theta_i'(\boldsymbol{\theta})=\int_{-\infty}^{\theta_i} p(\theta_i|\theta_{i+1},\cdots,\theta_m)\mathrm{d}\theta_i$ and then $p(\boldsymbol{\theta'})$ is a constant.

factor, penalising models with greater complexity [2]. Once suitable prior volumes have been fixed, the Hessian approximation to the hyperevidence is a computationally inexpensive means of comparing covariance functions.

The Hessian may also be used to provide error estimates for the hyperparameters; from (2.10) it can be seen that the inverse of the Hessian is the covariance matrix of the maximum hyperlikelihood estimator of the hyperparameters.

(b) Partial analytic maximisation

In general, covariance functions can be arbitrarily complicated, with large numbers of hyperparameters. Inevitably, simple covariance functions are the most prevalent in the literature. If there are a small number of hyperparameters, then even reducing the number of hyperparameters by one can have a great impact on the length of time taken to maximise the hyperlikelihood. In this section, we show how the hyperlikelihood for any covariance function, regardless of complexity, can be analytically maximised over an overall scale parameter, thereby reducing the number of remaining hyperparameters. We also generalise the expressions for the gradient and the Hessian found in Sec. 2(a) to this case.

Consider the following transformation of the covariance, $k(\boldsymbol{x_i}, \boldsymbol{x_j}) \to \sigma_f^2 k(\boldsymbol{x_i}, \boldsymbol{x_j})$; substituting this into the expression for the hyperlikelihood gives,

$$\ln P(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta}) = -\frac{1}{2\sigma_f^2}\boldsymbol{y}^{\mathrm{T}}\,\mathbf{K}^{-1}\,\boldsymbol{y} - \frac{1}{2}\ln(\det\mathbf{K}) - \frac{n}{2}\ln(2\pi\sigma_f^2)\,. \tag{2.14}$$

This function always has a unique maximum with respect to variations in σ_f^2 at the position

$$\hat{\sigma}_f^2 = \frac{1}{n}\boldsymbol{y}^{\mathrm{T}}\,\mathbf{K}^{-1}\,\boldsymbol{y}\,; \tag{2.15}$$

at this point the hyperlikelihood takes the value

$$\ln P_{\max}(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\vartheta}) = -\frac{n}{2}\ln\left(2\pi e\hat{\sigma}_f^2\right) - \frac{1}{2}\ln(\det\mathbf{K})\,. \tag{2.16}$$

Equation (2.16) is to be considered as a function of the remaining $m-1$ hyperparameters $\boldsymbol{\vartheta} = \{\boldsymbol{\theta} \setminus \sigma_f\}$. The peak evidence may now be found more easily by numerically maximising $\ln P_{\max}$ in (2.16) with respect to the remaining parameters $\boldsymbol{\vartheta}$. If a gradient-based algorithm is used, it is advantageous to have an analogous expression to (2.7) to give inexpensive derivatives. This can be found by differentiating (2.16) with respect to $\boldsymbol{\vartheta}$, making use of the results in (2.8),

$$\partial_{\boldsymbol{\vartheta}} \ln P_{\max}(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\vartheta}) = \frac{1}{2\hat{\sigma}_f^2}\boldsymbol{y}^{\mathrm{T}}\,\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}}\mathbf{K}\cdot\mathbf{K}^{-1}\,\boldsymbol{y} - \frac{1}{2}\,\mathrm{Tr}\left(\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}}\mathbf{K}\right)\,. \tag{2.17}$$

These are not the same as the derivatives in (2.7).

As well as maximising, we can also consider marginalising over σ_f [16]. As we are marginalising over a scale parameter we use the (improper) Jeffreys prior $P(\sigma_f) = c/\sigma_f$ [17]. The result is equal to the maximised form, up to a multiplicative constant,

$$P_{\mathrm{marg}}(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta}) = \int_0^\infty \mathrm{d}\sigma_f\,\frac{c}{\sigma_f}P(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta}) = \frac{c}{2}\left(\frac{2e}{n}\right)^{n/2}\Gamma\left(\frac{n}{2}\right)P_{\max}(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta})\,. \tag{2.18}$$

As before, once the peak hyperlikelihood has been found, the Hessian at the peak position can aid in model comparison. In this case, the Hessian should be calculated using the second derivatives of $\ln P_{\mathrm{marg}}$. However, we may instead differentiate $\ln P_{\max}$, as this differs only by a constant which will cancel when using the Hessian to compare two models. Differentiating (2.17)

with respect to $\boldsymbol{\vartheta}'$,[4]

$$\partial_{\boldsymbol{\vartheta}}\partial_{\boldsymbol{\vartheta}'}\ln P_{\mathrm{marg}}|_{\hat{\boldsymbol{\vartheta}}} \propto \frac{1}{2n\hat{\sigma}_f^4}\boldsymbol{y}^{\mathrm{T}}\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}}\mathbf{K}\cdot\mathbf{K}^{-1}\boldsymbol{y}\times\boldsymbol{y}^{\mathrm{T}}\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}'}\mathbf{K}\cdot\mathbf{K}^{-1}\boldsymbol{y}$$
$$-\frac{1}{2\hat{\sigma}_f^2}\boldsymbol{y}^{\mathrm{T}}\left[2\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}}\mathbf{K}\cdot\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}'}\mathbf{K}\cdot\mathbf{K}^{-1}-\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}}\partial_{\boldsymbol{\vartheta}'}\mathbf{K}\cdot\mathbf{K}^{-1}\right]\boldsymbol{y}$$
$$+\frac{1}{2}\operatorname{Tr}\left(\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}}\mathbf{K}\cdot\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}'}\mathbf{K}-\mathbf{K}^{-1}\cdot\partial_{\boldsymbol{\vartheta}}\partial_{\boldsymbol{\vartheta}'}\mathbf{K}\right). \tag{2.19}$$

Again, these are not the same as the derivatives in (2.9). These expressions for the gradient and the Hessian of the hyperlikelihood, maximised or marginalised over σ_f^2, share the same advantages as the analogous expressions in Sec. 2: they may be evaluated in $\mathcal{O}(n^2)$ time once the hyperlikelihood itself has been evaluated in $\mathcal{O}(n^3)$ time.

3. Numerical results

In order to perform model comparison calculations between competing covariance functions, we must first specify at least two different covariance functions. We choose the two functions in (3.1) and (3.2), where $(t,t')\equiv(x,x')$. These functions are both based on the periodic covariance function proposed by [2]. The first function k_1 is the product of a single periodic component with timescale T_1 and a simple compact-support polynomial covariance function [18] to describe any non-periodic component of the data. The choice of a compact-support covariance function is especially useful when working with large datasets; this is precisely the situation where the techniques described above are also designed to be of maximum benefit. The second function k_2 includes an additional periodic component with timescale T_2. In order to avoid double-counting in k_2, we impose the constraint $T_2\geq T_1$. Both covariance functions also include an uncorrelated noise term; we define this in such a way that σ_f remains an overall scale hyperparameter which can be maximised or marginalised over analytically as described in Sec. 3(b).

$$k_1(t,t')=\sigma_f^2 C\left(\frac{|t-t'|}{T_0}\right)\exp\left[-\frac{2}{l_1^2}\sin^2\left(\frac{\pi(t-t')}{T_1}\right)\right]+\sigma_f^2\sigma_n^2\delta_{tt'}\,, \tag{3.1}$$

$$k_2(t,t')=\sigma_f^2 C\left(\frac{|t-t'|}{T_0}\right)\exp\left[-\frac{2}{l_1^2}\sin^2\left(\frac{\pi(t-t')}{T_1}\right)-\frac{2}{l_2^2}\sin^2\left(\frac{\pi(t-t')}{T_2}\right)\right]+\sigma_f^2\sigma_n^2\delta_{tt'}\,, \tag{3.2}$$

$$C(\tau)=\begin{cases}(1-\tau)^5\dfrac{48\tau^2+15\tau+3}{3} & \tau<1\\ 0 & \tau>1\end{cases}. \tag{3.3}$$

The covariance functions are completely specified by the hyperparameters σ_f (overall scale), T_j ($j=0,1,2$; timescales), and l_j ($j=1,2$; smoothing parameters for the periodic components). The noise parameter σ_n could also be taken to be a hyperparameter; instead, for simplicity, we here take σ_n to be fixed. As σ_n appears in k multiplied by the overall scale, σ_f, fixing σ_n is roughly equivalent to specifying a fixed fractional error.

We want to perform model comparison using the Laplace approximation outlined previously. This technique requires reparametrising the covariance function such that the hyperpriors are flat. For the timescale hyperparameters, which are dimensionful, we choose to use the scale-invariant Jeffreys prior, $P(T_j)\propto 1/T_j$. This prior is improper if the range of T_j is $(0,\infty)$, so we restrict the range to $(\delta t,\Delta T)$, where δt and ΔT are respectively the smallest and largest separations between the sampling points. If there was a timescale in the problem outside of this range, we would be unable to resolve it from the data. We now seek a transformation $\phi_j\equiv\phi_j(T_j)$ to a new hyperparameter ϕ_j such that the prior is flat in this parameter, $P(\phi_j)=\mathrm{const}$. The conservation

[4]Here, $\hat{\sigma}_f$ retains its maximum $\ln P$ value from (2.15), although the new maximum $\ln P_{\mathrm{marg}}$ value has actually now shifted to $(\hat{\sigma}_f')^2=n\hat{\sigma}_f^2/(n-1)$ due to the effect of the hyperprior. For large data sets ($n\gg 1$) the difference between the two is negligible (the hyperprior becomes uninformative as it is overwhelmed by the hyperlikelihood).

of probability gives a differential equation relating the two

$$P(T_j)\mathrm{d}T_j = P(\phi_j)\mathrm{d}\phi_j \quad \Rightarrow \quad T_j = \exp(\phi_j/A_j)\,, \quad \{j=0,1,2\}\,, \tag{3.4}$$

where the A_js are constants which we can set equal to 1. The range of these new hyperparameters is $\phi_j \in (\ln(\delta t), \ln(\Delta T))$ and $P(\phi_j) = 1/\ln(\Delta T/\delta t)$.

For the smoothness parameters l_j we choose to use log-normal priors, $P(l_j) = \exp\left[-(\mu - \log l_j)^2/(2\sigma_l^2)\right]/\sqrt{2\pi\sigma_l^2}$, with mean $\mu = 1$ and variance $\sigma_l^2 = 4$. As before, we seek a transformation to some new hyperparameters ξ_j in which the prior is flat. The desired transformation is given by

$$l_j = \exp\left[\mu + \sqrt{2}\sigma_l\,\mathrm{erf}^{-1}(2\xi_j)\right]\,, \quad \{j=1,2\}\,, \tag{3.5}$$

where $\xi_j \in (-0.5, 0.5)$.

(a) Synthetic data

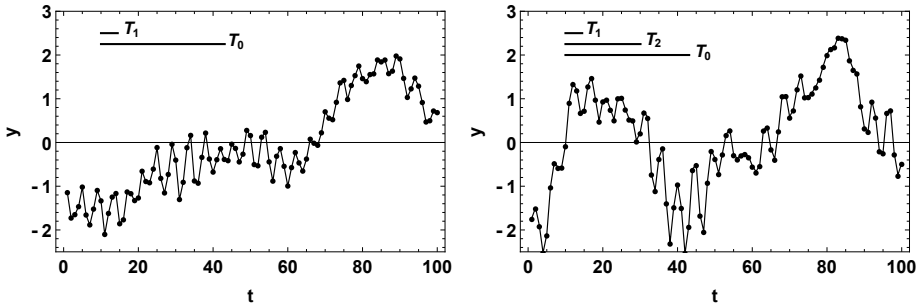

Figure 1. Realisations of the GPs $k_1(t,t')$ and $k_2(t,t')$ from (3.1) and (3.2) for values of $t = 1, 2, 3 \ldots, 100$ are shown in the left and right hand panels respectively. The horizontal black lines indicate the lengthscales associated with the different terms in the covariance functions. The hyperparameters for k_1 where chosen to be $\sigma_f = 1$, $\phi_0 = 3.5$, $\phi_1 = 1.5$, and $x_1 = 0$. The hyperparameters for k_2 where chosen to be the same as for k_1 and $\phi_2 = 3$ and $x_2 = 0$. In both cases the noise was fixed to $\sigma_n = 10^{-2}$

Shown in Fig. 1 are realisations of GPs with covariance functions k_1 and k_2.[5] In order to perform test model comparison calculations, a realisation of the k_2 GP with n points was drawn and analysed using both the k_1 and k_2 covariance functions. For each covariance the peak hyperlikelihood was found by numerically maximising (2.14) using a conjugate gradient method, making use of the gradient in (2.17). The hyperevidence was estimated using (2.13) and the expression for the Hessian in (2.19); the results are summarised in Tab. 1. To verify the accuracy of this estimate, the hyperevidence was also integrated numerically using MULTINEST, [19–21] which implements a nested sampling algorithm [22]. This was repeated for three different values of n (in the case $n = 100$ the synthetic data is plotted in the right-hand panel of Fig. 1), and the results are also summarised in Tab. 1.

From Tab. 1 it can be seen that as n is increased, the Bayes factors increasingly favour the more complicated covariance function (and in this case the correct covariance function from which the data was drawn). In almost all cases, the Laplace approximation gives a value $\ln \mathcal{Z}_{\mathrm{est}}$ which is in agreement at better than 2σ with the numerically integrated value $\ln \mathcal{Z}_{\mathrm{num}}$. There is one exception which is highlighted in bold; this occurs for the most complicated covariance function (with the largest number of hyperparameters) and when the number of data points is smallest. In this situation, it would be expected that the posterior distribution on the

[5]The code, optimised for use on a GPU, and the synthetic used to produce these numerical results is available at http://hdl.handle.net/10283/1924.

n	$\ln \mathcal{Z}^{k_1}_{\rm est}$	$\ln \mathcal{Z}^{k_1}_{\rm num}$	$\ln \mathcal{Z}^{k_2}_{\rm est}$	$\ln \mathcal{Z}^{k_2}_{\rm num}$	$\ln \mathcal{B}_{\rm est}$	$\ln \mathcal{B}_{\rm num}$
30	-17.77	-17.87 ± 0.08	$\mathbf{-18.82}$	-17.73 ± 0.09	-1.05	0.14 ± 0.12
100	-20.17	-20.17 ± 0.10	-19.22	-19.22 ± 0.11	0.95	0.95 ± 0.15
300	-49.94	-50.12 ± 0.11	-40.21	-40.36 ± 0.13	9.73	9.76 ± 0.17

Table 1. A summary of the results of the analysis of synthetic data for three different-sized data sets. The first set of two columns is for a data set drawn from the k_2 covariance function and analysed with the k_1 covariance function. The first column is the estimated hyperevidence using the Laplace approximation where $\mathcal{Z}$ is as given in Eq. (2.13), while the second is the numerically calculated hyperevidence. The second set of two columns shows results for the same data, but analysed with the k_2 covariance function. The final pair of columns shows the log Bayes factor, $\ln \mathcal{B} \equiv \ln \mathcal{Z}^{k_2} - \ln \mathcal{Z}^{k_1}$, calculated using the approximate and numerical values for the hyperevidence.

hyperparameters may be highly multimodal and/or exhibit strong degeneracies (both of these expectations were confirmed by examining the posterior distribution on the hyperparameters returned by MULTINEST). This exceptional case serves to highlight situations in which the Laplace approximation should not be trusted. The MULTINEST posteriors in all other cases were verified to be well approximated by a single Gaussian mode. Fig. 2 shows the posterior distribution for the parameters of k_2 obtained from the largest ($n = 300$) synthetic data set.

Our method of model comparison is proposed as a faster alternative to model comparison using numerically evaluated Bayes factors. Simply comparing the peak hyperlikelihood (marginal likelihood) values would also give a measure of the goodness-of-fit, but this tends to favour more complex models and incurs the risk of overfitting. More sophisticated methods of model selection exist in the literature (see [23, 24] and references within), e.g., the comparison of models based on estimated predictive criteria [25–27], or the construction of a larger reference model and the subsequent selection of a simpler submodel with similar predictions [28, 29]. A detailed numerical comparison to these methods is left for future investigation.

The $\ln \mathcal{Z}_{\rm num}$ values in Tab. 1, evaluated using MULTINEST, required between 20,000 and 50,000 likelihood evaluations. The maximisation routines typically took fewer than 100 likelihood evaluations to find the peak, and then one additional evaluation to calculate the Hessian and hence $\ln \mathcal{Z}_{\rm est}$. In order to guard against the possibility of the maximisation routines becoming trapped in local maxima, as opposed to the global maximum, the algorithm was run multiple times from randomly selected starting positions. The typical number of runs required to find the global maximum was ~ 10. After these duplicate runs are accounted for, the speed-up factor in calculating $\ln \mathcal{Z}_{\rm est}$ compared to $\ln \mathcal{Z}_{\rm num}$ was between 20 and 50 in all cases.

(b) Tidal data from Woods Hole

In order to illustrate the effectiveness of the techniques described above on real data, we consider several tidal data sets of different sizes from Woods Hole, MA [30].[6] We consider the mean sea level offset recorded between 3 January and 15 June 2014, six lunar months, sampled at two-hour intervals (giving $n = 1968$ data points). This is plotted in Fig. 3. We also consider a smaller subset of the data (the first lunar month), with $n = 328$ data points.[7]

We interpolate the data using the two covariance functions in (3.1) and (3.2); these functions are well suited to the data, as we expect the sea level to contain harmonics of the various timescales associated with the daily, monthly and yearly cycles of the tides. For simplicity, we fix $\sigma_n = 10^{-2}$, which is the typical fractional error in the sea-level measurements. As in Sec. 3(a) we reparametrise the covariance functions so that the T_j have Jeffreys priors, and the smoothness parameters have log-normal priors. We use a conjugate gradient maximisation algorithm with

[6] Data from the National Oceanic and Atmospheric Administration, http://tidesandcurrents.noaa.gov/waterlevels.html (accessed August 2015).

[7] This data set is regularly sampled in time, and therefore the covariance matrix will be a Toeplitz matrix. This structure could be exploited to accelerate the inversion of the covariance matrix; we choose not to use this here so that our code can be applied to irregularly sampled data in arbitrary dimensions.

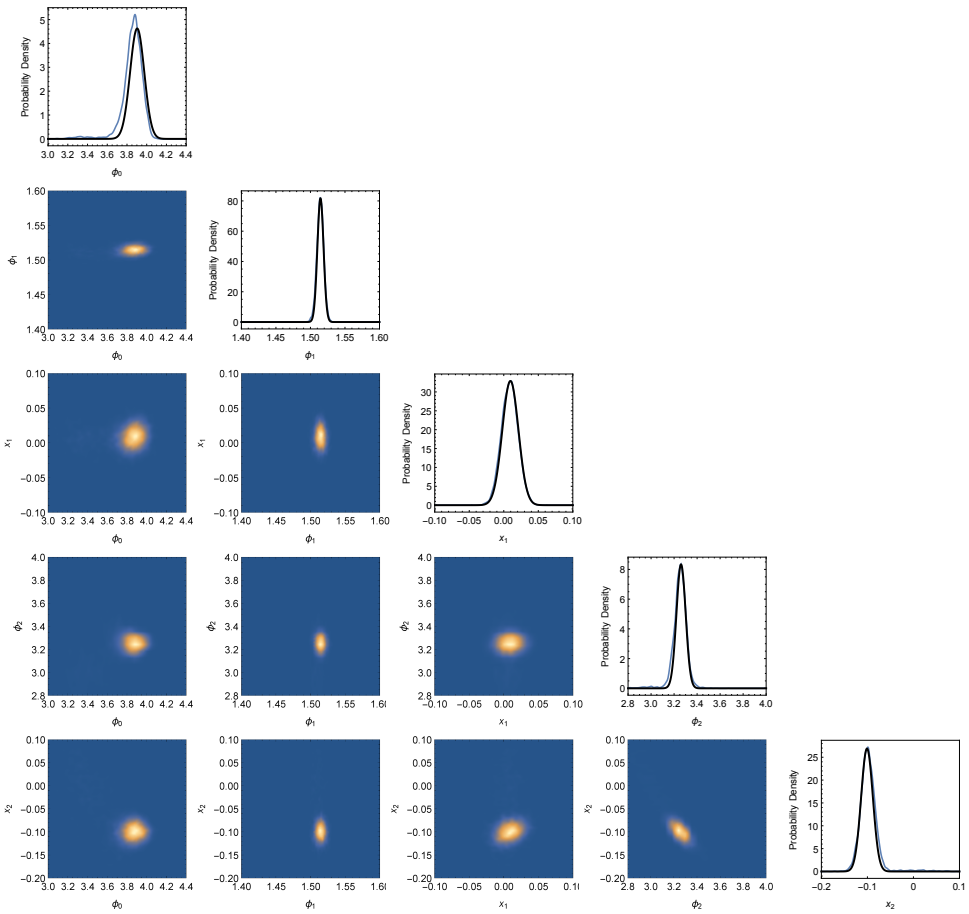

Figure 2. The one and two dimensional marginalised posterior distributions on the hyperparameters of the k_2 covariance function obtained from the largest ($n = 300$) synthetic data set. The posterior is well approximated as a normal distribution. Shown in the black curves in the one dimensional marginalised posterior distributions along the diagonal are the normal approximations obtained by using the techniques described in Sec. 2 to maximise and find the Hessian. Using the Hessian to approximate the integral of this distribution (the hyperevidence) leads to an error of $\sim 10\%$, see Tab. 1.

(2.17) and the Hessian in (2.19) to evaluate the volume in (2.13) and perform model comparison between the two covariance functions.

For the smaller dataset, we find the timescale $T_1 = (12.8 \pm 0.2)$ hours with k_1, which corresponds to the two main tides per day. With k_2 we find the timescales $T_1 = (12.44 \pm 0.07)$ hours and $T_2 = (24.3 \pm 1.0)$ hours; the second timescale corresponds to the height difference between the first and second tides of the day. The two-timescale model is highly favoured with a log Bayes factor of 57.8.

For the larger dataset, we find $T_1 = (12.80 \pm 0.11)$ hours with k_1, and $T_1 = (12.40 \pm 0.03)$ hours and $T_2 = (23.3 \pm 0.3)$ hours with k_2. In all cases, the (squared) errors are estimated using the diagonal components of the inverse Hessian; it can be seen that the timescales are more precisely measured for the larger dataset, as expected. The two-timescale model is even more conclusively favoured for the larger dataset, with a log Bayes factor of 538. We also find a number of subsidiary hyperlikelihood peaks associated with other timescales in the data, but all subsidiary peaks are

strongly suppressed relative to the global peak (by at least $\Delta \ln P$ of ~ 100) and so we expect our Bayes factor estimates to be robust.

Sea-level data is known to contain a large number of different frequencies, which necessitates the use of harmonic analysis in tidal modelling; the number of constituents included in tide prediction calculations has increased from tens [31] to thousands [32] over the past century. Clearly any k_2-like covariance function with < 10 timescales is simplistic, but the construction of a more detailed tidal model is beyond the scope of this paper. [8]

The number of evaluations of (2.13) needed to obtain these results was comparable to the numbers for the synthetic data discussed in Sec. 3(a). However, each evaluation here was more expensive (~ 10 s) due to the size of the data set. Based on the speed-ups found in Sec. 3(a), it would be expected that MULTINEST would take up to ~ 1 week to calculate the Bayes factor.

Shown in the inset plot in Fig. 3 are the two interpolants from k_1 and k_2 for the larger dataset, which both perform equally well on the timescale of one week. These interpolants, which are the result of the regression analysis, may be used to estimate the tidal height at a time where a measurement is not available.

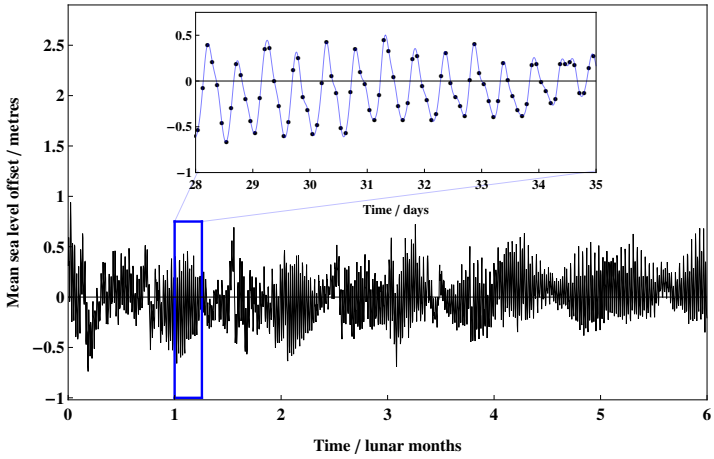

Figure 3. Shown in the main figure are six lunar months of tidal height data (black), from which the lunar tidal cycle can be discerned. Shown in the inset plot are several days of the tidal data (black points), from which the daily cycles can be clearly seen. Overlaid in the inset plot are both GP interpolants (blue), which are identical on this timescale.

4. Summary

We have described some simple ways in which the computationally expensive training stage of implementing GPR can be accelerated. The analytic maximisation of the hyperlikelihood over a single scale hyperparameter of the covariance function aids in speeding up the maximisation of the hyperlikelihood by reducing the dimensionality of the problem; the advantages of this will be most keenly felt in (common) problems where relatively simple covariance functions are used. Meanwhile, the analytic evaluation of the Hessian matrix, either in the manner of (2.9) or (2.19), aids in speeding up the process of model comparison between different types of covariance function. We have successfully demonstrated these techniques on a synthetic data set where the data was drawn from one of two covariance functions under consideration. In the case of the synthetic data the size of the data set was limited to less than 300 points, so that the results

[8] We have conducted preliminary investigations of a three-timescale model. The hyperlikelihood surface for this covariance function is more structured and non-Gaussian than for k_2 and k_1. Estimates of the Bayes factors indicate that the inclusion of additional timescales is favoured, as expected based on the known large number of modes present.

could be verified by using the MULTINEST algorithm to numerically sample and integrate the posteriors. We also demonstrated the techniques by applying them to a larger real data set of mean sea level measurements, where the full MULTINEST calculation would have taken too long to perform. It is to be hoped that these techniques will aid in the wider application of GP methods to larger data sets.

Data accessibility The original source of the tidal data is the National Oceanic and Atmospheric Administration (http://tidesandcurrents.noaa.gov/waterlevels.html, accessed August 2015). The subset of data used here, as well as our synthetic data sets and our code is available at http://hdl.handle.net/10283/1924.

Funding statement CJM and CPLB are supported by the STFC. AJKC's work is supported by the Cambridge Commonwealth, European and International Trust. JRG's work is supported by the Royal Society.

Author contributions CJM devised the study, jointly performed the analysis and wrote the manuscript. AJKC jointly performed the analysis and wrote the manuscript. CPLB helped with the analysis, and the writing and editing of the manuscript. JRG helped devise the study and edited the manuscript. All authors gave final approval for the publication.

Conflict of interest We have no competing interests.

References

1 C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
2 D. J. C. MacKay, *Information Theory, Inference and Learning Algorithms* (Cambridge University Press, Cambridge, 2003).
3 D. Barry, The Annals of Statistics **14**, 934 (1986).
4 G. Wahba, Journal of the Royal Statistical Society. Series B (Methodological) **40**, 364 (1978).
5 A. O'Hagan and J. F. C. Kingman, Journal of the Royal Statistical Society. Series B (Methodological) **40**, 1 (1978).
6 A. J. Smola and P. L. Bartlett, in *Advances in Neural Information Processing Systems*, edited by T. K. Leen, T. Dietterich, and V. Tresp (MIT Press, Cambridge, MA, 2001), vol. 13, pp. 619–625.
7 J. Quiñonero-Candela and C. E. Rasmussen, Journal of Machine Learning Research **6**, 1939 (2005).
8 N. Cressie and G. Johannesson, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **70**, 209 (2008).
9 A. Banerjee, D. B. Dunson, and S. T. Tokdar, Biometrika **100**, 75 (2013), 1106.5779.
10 C. J. Moore, C. P. L. Berry, A. J. K. Chua, and J. R. Gair, Physical Review D **93**, 064001 (2016), 1509.04066.
11 C. J. Moore and J. R. Gair, Physical Review Letters **113**, 251101 (2014), 1412.3657.
12 D. J. C. MacKay, Neural Computation **11**, 1035 (1999).
13 E. Snelson and Z. Ghahramani, in *Advances in Neural Information Processing Systems*, edited by Y. Weiss, B. Schölkopf, and J. C. Platt (MIT Press, Cambridge, MA, 2006), vol. 18, pp. 1257–1264.
14 J. Quiñonero-Candela, C. E. Rasmussen, and C. K. I. Williams, in *Large-scale Kernel Machines*, edited by L. Bottou, O. Chapelle, D. DeCoste, and J. Weston (MIT Press, Cambridge, MA, 2007), chap. 9, pp. 203–223.
15 M. Blum and M. Riedmiller, in *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* (2013).
16 D. J. C. MacKay, in *Maximum Entropy and Bayesian Methods*, edited by G. R. Heidbreder (Springer Netherlands, Dordrecht, 1996), vol. 62 of *Fundamental Theories of Physics*, pp. 43–59.

17 H. Jeffreys, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **186**, 453 (1946).

18 H. Wendland, *Scattered Data Approximation (Cambridge Monographs on Applied and Computational Mathematics)* (Cambridge University Press, 2005).

19 F. Feroz, M. P. Hobson, and M. Bridges, Monthly Notices of the the Royal Astronomical Society **398**, 1601 (2009), 0809.3437.

20 F. Feroz and M. P. Hobson, Monthly Notices of the Royal Astronomical Society **384**, 449 (2008), 0704.3704.

21 F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt, *Importance Nested Sampling and the MultiNest Algorithm* (2013), 1306.2144.

22 J. Skilling, Bayesian Analysis **1**, 833 (2006).

23 R. B. O'Hara and M. J. Sillanpää, Bayesian Analysis **4**, 85 (2009).

24 J. Piironen and A. Vehtari, *Comparison of Bayesian predictive methods for model selection* (2015), 1503.08650.

25 S. Geisser and W. F. Eddy, Journal of the American Statistical Association **74**, 153 (1979).

26 S. Watanabe, *Algebraic geometry and statistical learning theory*, Cambridge Monographs on Applied and Computational Mathematics (Cambridge University Press, Cambridge, 2009).

27 A. Gelman, J. Hwang, and A. Vehtari, Statistics and Computing **24**, 997 (2013), 1307.5928.

28 D. V. Lindley, Journal of the Royal Statistical Society. Series B (Methodological) **30**, 31 (1968).

29 A. San Martini and F. Spezzaferri, Journal of the Royal Statistical Society. Series B (Methodological) **46**, 296 (1984).

30 W. Scherer, W. M. Stoney, T. N. Mero, M. O'Hargan, W. M. Gibson, J. R. Hubbard, M. I. Weiss, O. Varmer, B. Via, D. M. Frilot, et al., Tech. Rep. NOAA Special Publication NOS CO-OPS 1, National Oceanic and Atmospheric Administration, Silver Spring, Maryland (2001).

31 E. Lisitzin, *Sea-level changes* (Elsevier, Oxford, 1974).

32 S. Casotto and F. Biscani, in *AAS/Division of Dynamical Astronomy Meeting #35* (2004), vol. 36 of *Bulletin of the American Astronomical Society*, p. 862.